Evolving an expert checkers player without using human expertise

Kumar Chellapilla
Natural Selection, Inc.
3333 N. Torrey Pines Ct., Suite 200
La Jolla, CA 92037
kumar@natural-selection.com

Humans as a Benchmark

- Since the advent of the modern digital computer we've tried to generate machines as intelligent as humans
- Primary goal of early artificial intelligence
  - General problem solver
- What is intelligence?
  - Central problem: No well-accepted definition of intelligence, let alone "artificial intelligence"
- Surrogate challenges in the form of Turing Tests

Turing Test

- Turing (1950) replaced the question of “Can a machine think?” to “Can a machine fool an interrogator as well as a human?”
- The test involves a man trying to convince an interrogator than he is a woman
- Machine “passes the test” if it can fool the interrogator as often as the man

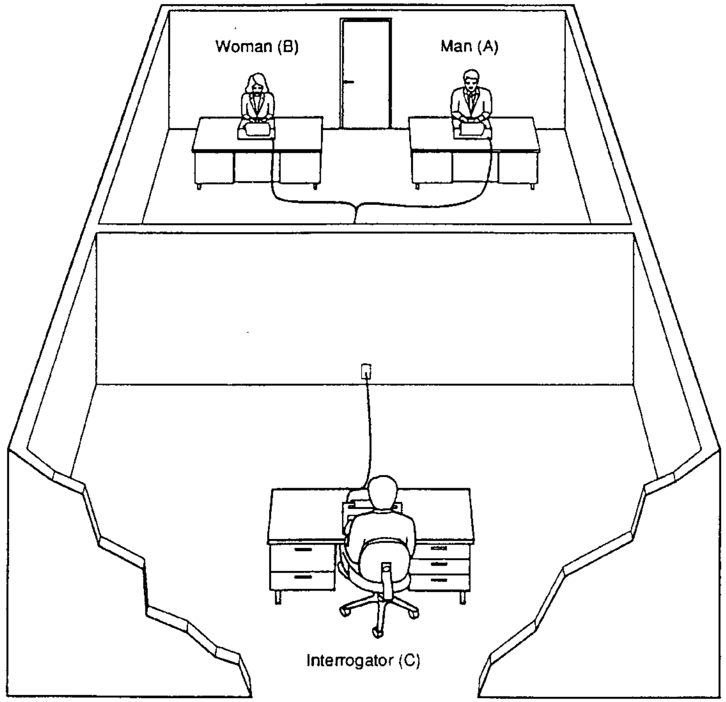

Figure 1-1 The Turing test. An interrogator (C) questions both a man (A) and a woman (B) and attempts to determine which is the woman.

Turing Test $\rightarrow$ Intelligence?

- Turing (1950) never claimed that a machine that passes the test would be "intelligent"
  - "Too meaningless to deserve discussion"
- The Turing Test is no more of a test for thinking machines than it is a test for femininity
  - If a man can fool an interrogator into believing he is a woman as often as a woman can convince the interrogator, is the man a woman?
- What were the consequences of focusing on this test?

Consequences

- Impossible to envision passing the test in the 1960s based purely on computer speed
- Narrow the focus
  - Try games
- Emphasize applications, emphasize humans
- Ask experts how they do things
- Never mind what intelligence is what sort of intelligence we are trying to generate
  - Minsky: Intelligence means the ability to solve hard problems
- The Turing Test led to the death of AI

Artificial Intelligence

- For an organism (system) to be intelligent, it must make decisions
- A decision arises when available resources must be allocated
  - Must face a range of decisions, otherwise there's really no decision at all
- Decision making requires a goal
- Intelligence may be defined as "the ability of a system to adapt its behavior to meet its goals in a range of environments"

Adaptive Behavior $\rightarrow$ Evolutionary Computation

- Adaptation is fundamentally an evolutionary process whether it occurs in phyletic, ontogenetic, or sociogenetic systems
- Unit of mutability and reservoir of stored knowledge
- The mechanisms for change and memory differ but the behavioral effects are notably similar
- If we really want to talk about intelligent machines we have to talk about machines that learn and adapt to meet goals based on experience
  - MACHINES THAT EVOLVE

Where Does That Leave The Turing Test?

- Comparing the quality of evolved behavior to that of humans is a perfectly reasonable endeavor
- It measures the quality of behavior not the “intelligence” of the behavior or of the underlying system
- Quality of behavior is important!
- So then: How well does evolutionary computation stack up against humans?

Early Evolutionary Computation

- Fogel (1964; Fogel et al. 1966, p. 39) compared predictive capability of evolution of finite-state machines to humans on environments created by Merrill Flood
- Comparison: Equivalent on all but three instances, where the evolutionary algorithm was shown to be superior

Table 3.2
A Comparison of the Predictive Ability of the Evolutionary Program to That of Human Subjects—(M. M. Flood)

Initial Sequence	Subject	Trial	Subjects' Score	Evolutionary Program Score
5044014 . . .	*A*	1	0.425	0.460
2755776 . . .	*A*	2	0.445	0.515
0644114 . . .	*B*	1	0.435	0.490
2534756 . . .	*B*	2	0.455	0.435
5744204 . . .	*C*	1	0.445	0.505
2754356 . . .	*C*	2	0.420	0.510
0044304 . . .	*D*	1	0.435	0.560
6534356 . . .	*D*	2	0.545	0.505
0444014 . . .	*E*	1	0.310	0.565
2715756 . . .	*E*	2	0.460	0.450
0144304 . . .	*F*	1	0.495	0.465
2515356 . . .	*F*	2	0.545	0.520
6574356 . . .	*G*	1	0.500	0.495
5444104 . . .	*G*	2	0.570	0.495

More Comparisons

- Evolutionary algorithm used to predict symbols from Wolin (Human Factors, 1963)
- Based on different recall lengths, the evolutionary algorithm could
- Outperform the humans (college students)

Table 3.3

A Comparison of the Predictive Ability of the Evolutionary Program to that of Human Subjects (B. R. Wolin)

Initial Sequence	Designation of Environment	Number of Subjects	Ideal Score	Evolutionary Score (with Growth)	Subjects' Average Score	Evolutionary Score (10-Symbol Recall)
521212 . . .	TR-1	6	0.75	0.730	0.690	0.639
523222 . . .	TR-3	6	0.75	0.720	0.695	0.630
4000524 . . .	TR-13	6	1.0	0.730	0.700	0.575
05050524 . . .	TR-U1A	5	0.875	0.865	0.830	0.725
130271 . . .	TR-U3	7	0.92	0.860	0.850	0.725
64125 . . .	WORD	24	0.934	0.845	0.760	0.540

Evolutionary Fluid Mechanics

- Rechenberg, Schwefel, and Bienert at Technical University of Berlin evolved hardware designs that exceed human design

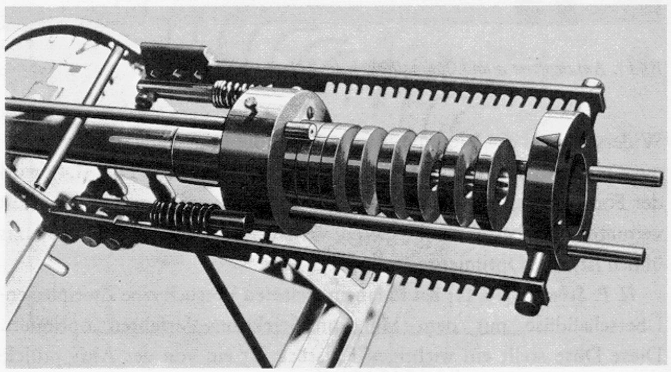

Bild 8. Versuchsaufbau – segmentierte Düse

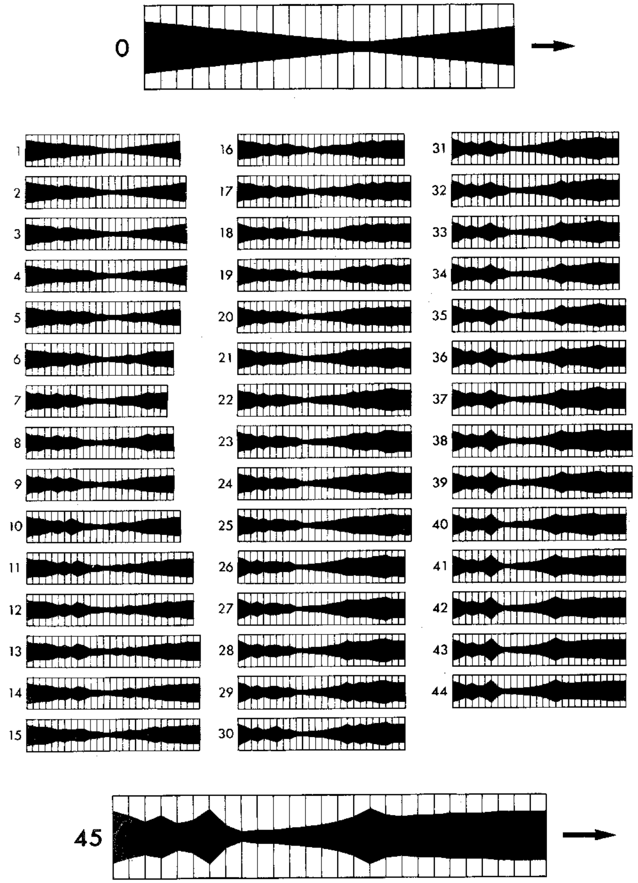

Bild 9. Entwicklung einer Zweiphasen-Überschalldüse von der Anfangsform 0 zur Optimalform 45

Evolving Strategies to Games

- Reed, Toombs, and Barricelli (1967) evolved strategies for a simplified game of poker
- High and low cards; High, low, and pass bets
- Strategies defined by probabilities, and included self-adaptation

- “In all three experiments, the high hand betting probabilities were nearly optimized in less than 200 generations and approached the optimum values … calculated by von Neumann’s game theory. The low hand betting probabilities, which are less important for the quality of the game, were not optimized yet and still presented considerable differences in different patterns. The quality of the game was fully competitive with average human players uninformed about game theory.”

More Recent Efforts in Evolutionary Computation for Gaming

- Axelrod (1987) evolved strategies in the iterated prisoner's dilemma that were competitive with strategies that did well in his previous tournaments
- Fogel (1993) evolved neural network strategies in tic-tac-toe that performed well against an "expert system"
- Moriarty and Mikkulainen (1995) evolved a neural network to play "above-average" Othello
- Pollack and Blair (1998) evolved a neural network to play backgammon
- What's missing are direct comparisons to humans

The Game of Checkers

- 8x8 board with red and black squares
- Two players (Red & White)
- 12 pieces (checkers) for each player
- Diagonal Moves
- Jumps are forced
- Checkers and kings
- Win, lose, and draw

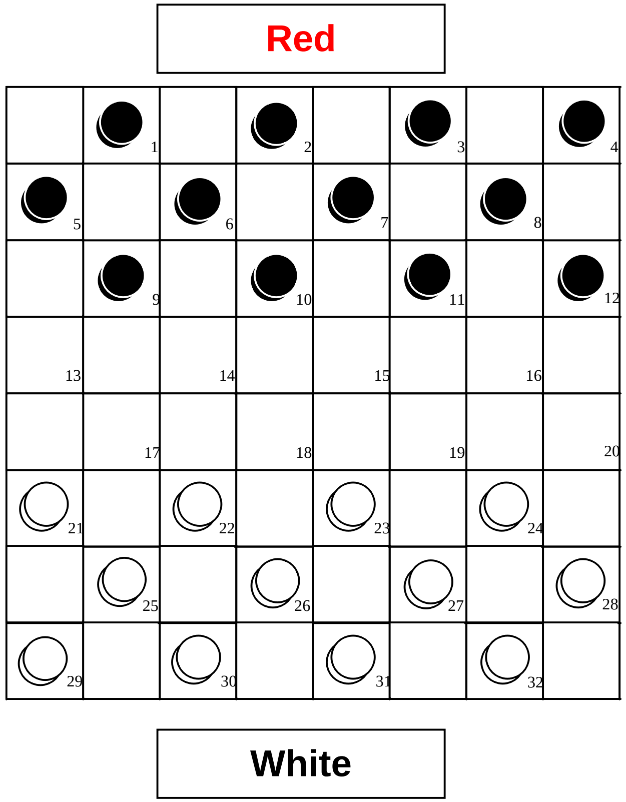

Computer Checkers

- Samuel's first checkers program
- World Man-Machine Checkers Championship
  - Chinook defeated Marion Tinsley (human), the world checkers champion and won the championship
- Chinook
  - Incorporated a linear polynomial as a board evaluator
  - All "items" of knowledge were preprogrammed, opening book, and all 8-piece endgame database (440 billion stored positions)
  - Did not use any learning
- Programmed human expertise to beat human expertise

Evolving Strategies for Checkers

- 32x1 board vector
- Entries $\{-K, -1, 0, 1, K\}$
- Players pieces positive
- Opponents pieces negative
- Each player consisted of
  - A neural network board evaluator
  - A unique king value K
  - The NN and K are evolvable
- Minimax search
  - 4-ply for training and 6-ply for playing against humans

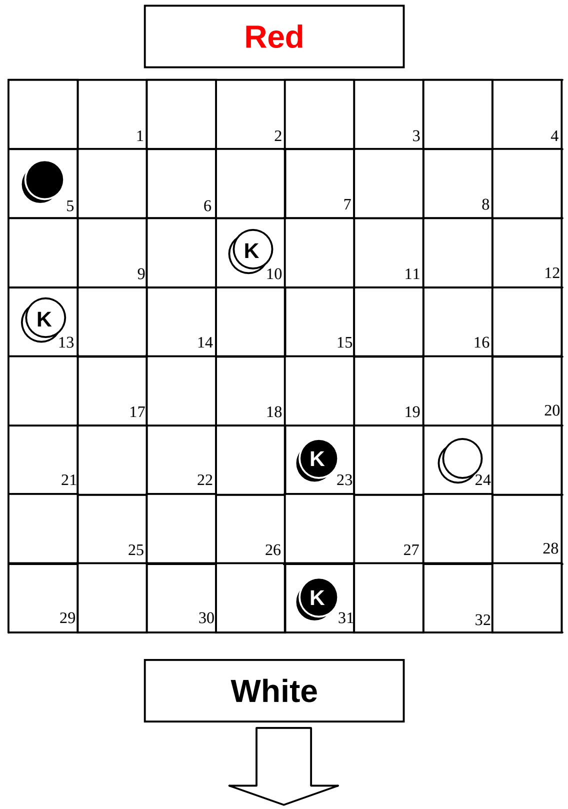

[0 … 0 1 0 … 0 -1.5 0 0 -1.5 0 … 1.5 -1 … 0 1.5 0]
5 10 13 23 24 31
King Value = 1.5

Neural Network I Architecture

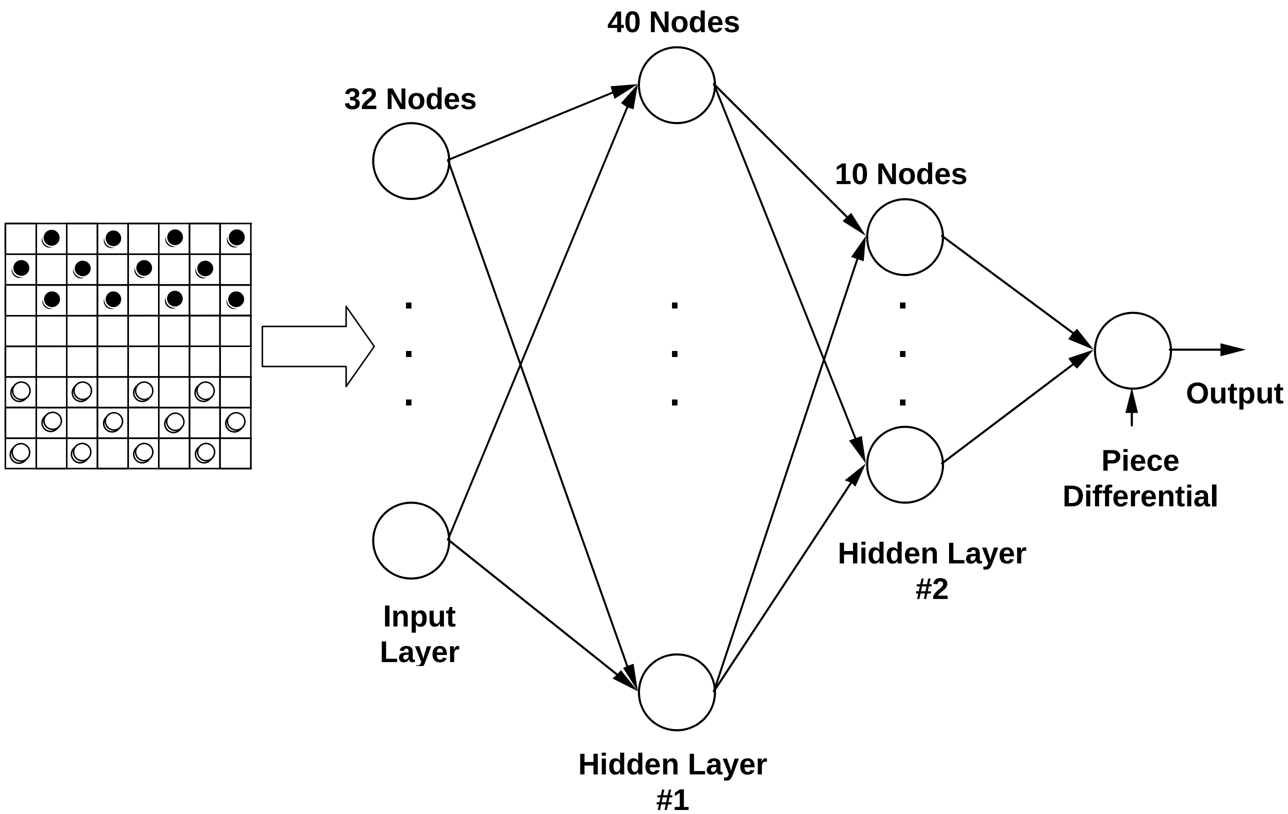

- The closer the NN output was to 1.0 the better the move
- The pieces changed sign when move alternated between players

Evolving Checkers Players

- Neural network weight update

$$\sigma'_i(j) = \sigma_i(j)\exp(\tau N_j(0,1)), \qquad j = 1, \ldots, N_w \text{ and } \tau = \left(\sqrt{2\sqrt{N_w}}\right)^{-1}$$

$$w'_i(j) = w_i(j) + \sigma'_i(j) N_j(0,1), \qquad j = 1, \ldots, N_w$$

- King value update

$$K' = K + 0.1U, \qquad \text{where } U \in \{-1,0,1\}$$

  – K was limited to [1.0,3.0]
- Tournament
  – Each player (parents and offspring) played one checkers game with five randomly selected opponents from the population
  – Win = 1 points, draw = 0 points, and loss = –2 points
  – Games were limited to a maximum of 100 moves

Evolution

- 0. Initialization
  - 15 parents with NN weights uniformly sampled from [–0.2,0.2]
- 1. Offspring generation
  - Each parent generated one offspring
- 2. Tournament
  - All 30 players competed with 5 randomly selected players from the population
- 3. Selection
  - 15 players with the greatest total points were retained as parents for the next generation
- 4. Loop back to step 1.
- Evolution was conducted for 100 generations

Evaluation Against Human Players

- Best player at generation 10 defeated the authors (novice checkers players)
- Best player at generation 100 was evaluated over 100 games against rated human players at the internet gaming site: www.zone.com
- USCF checkers rating on the zone
  - Starts out at 1600 and follows:

$$\text{Outcome} \in \{\ 1(\text{win}),\ 0.5(\text{draw}),\ 0(\text{loss})\ \}$$

$$W = \frac{1}{1 + 10^{0.0025\left(R_{opp} - R_{Old}\right)}}$$

$$R_{new} = R_{old} + 32(\text{Outcome} - W)$$

Results: Game Outcomes

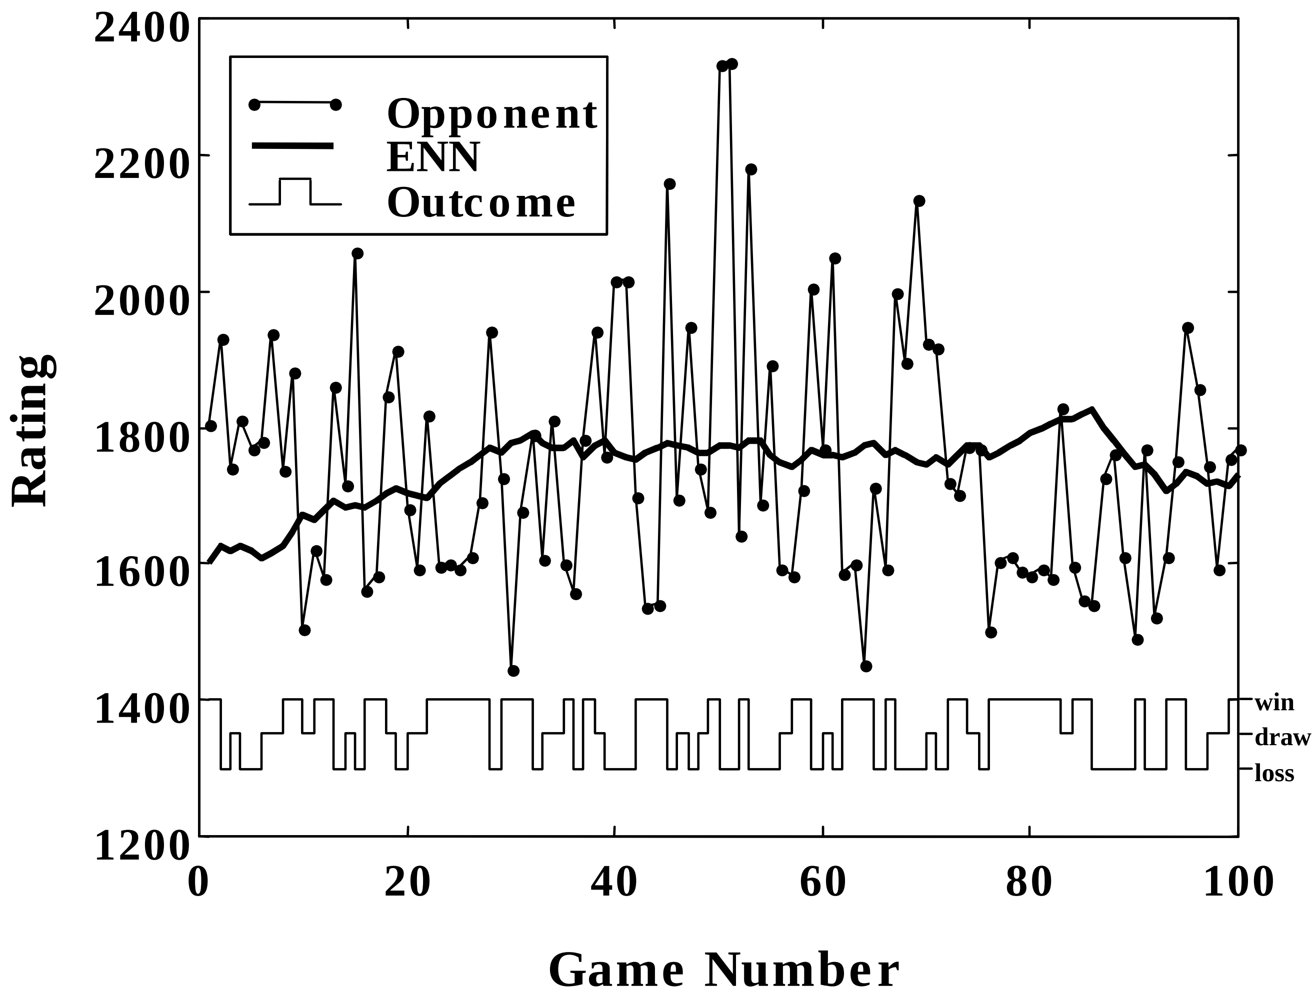

Results: Game Outcomes

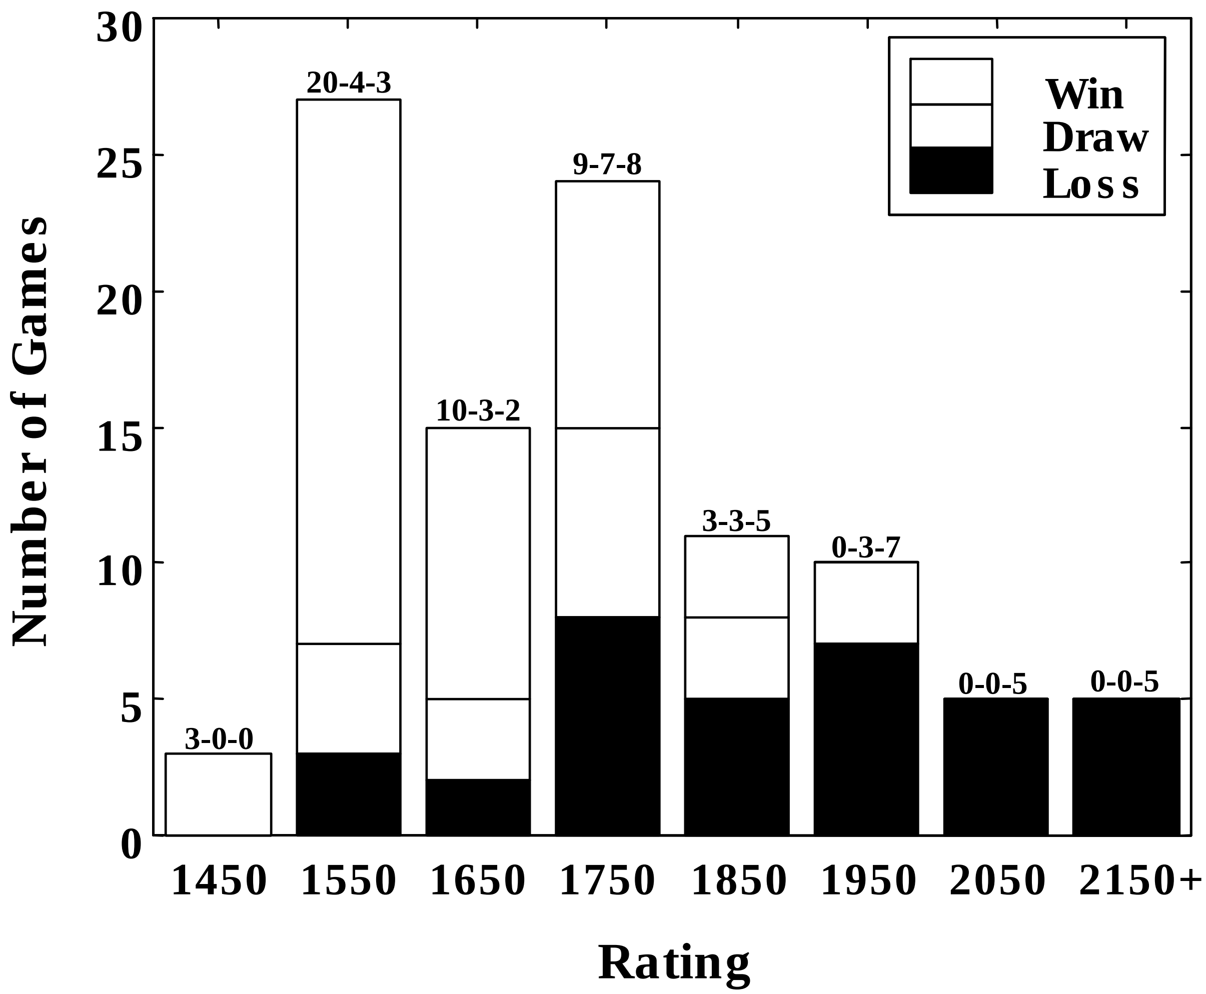

Game against Human rated 1946: Draw

Human	ENN
1.R:11-16,	W:23-19;
2.R:16-23(f),	W:26-19
3.R:8-11,	W:19-15
4.R:11-18,	W:22-15(f)
5.R:10-19(f),	W:24-15(f)
6.R:7-10,	W:27-24
7.R:10-19(f),	W:24-15(f)
8.R:6-10,	W:15-6(f)
9.R:1-10(f),	W:25-22
10.R:9-14,	W:30-26
11.R:3-7,	W:22-17
12.R:4-8,	W:26-23
13.R:8-11,	W:28-24
14.R:11-15,	W:32-28
15.R:7-11,	W:29-25
16.R:15-18,	W:23-19
17.R:18-23,	W:24-20
18.R:5-9,	W:17-13
19.R:14-18,	W:13-6(f)
20.R:2-9(f),	W:21-17
21.R:9-13,	W:19-15
22.R:10-19	

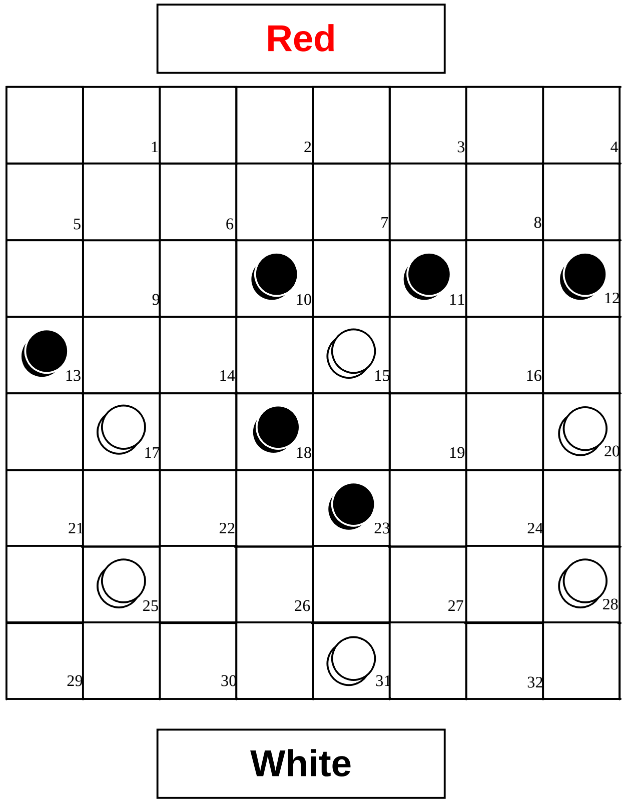

[take the piece on 15; frees piece on 17, Was this a mistake? Should have double jumped to get king 13-22-29?]

Game against Human rated 1946: Draw

Human	ENN
21.R:9-13,	W:19-15
22.R:10-19,	W:17-14
23.R:13-17,	W:25-21
24.R:17-22,	W:14-9
25.R:11-15,	W:9-6
26.R:22-26,	W:31-22(f)
27.R:18-25(f),	W:6-1
28.R:15-18,	W:1-6
29.R:23-27,	W:6-10
30.R:18-22,	W:10-15
31.R:19-23,	W:20-16
32.R:12-19(f),	W:15-24-31(f)
33.R:25-29,	W:21-17
34.R:22-25,	W:17-13
35.R:25-30,	W:13-9
36.R:30-25,	W:9-5
37.R:23-26,	W:31-22(f)
38.R:25-18(f),	W:5-1
39.R:29-25,	W:1-5
40.R:18-14,	W:28-24
41.R:25-22,	W:24-19
42.R:22-18,	W:19-16
43.R:18-15,	W:16-12

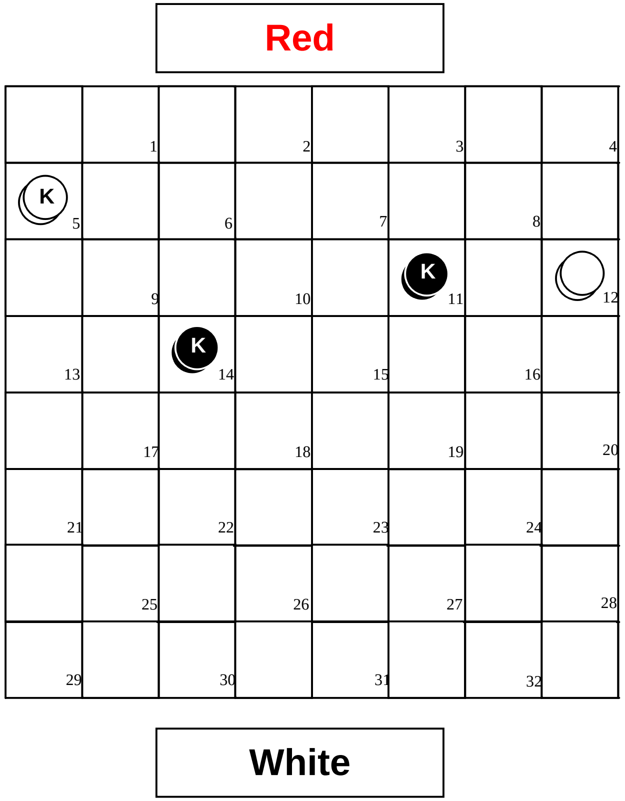

[After 10 more moves, game ends with red offering a draw]

Game against Human rated 1771: Win

ENN	Human
1.R:9-13,	W:22-18
2.R:11-15	W:18-11(f)
3.R:7-16	W:25-22
4.R:5-9	W:22-18
5.R:3-7	W:29-25
6.R:1-5	W:25-22
7.R:16-19	W:23-16
8.R:12-19(f)	W:24-15(f)
9.R:10-19(f)	W:27-24
10.R:7-11	W:24-15(f)
11.R:9-14	W:18-9(f)
12.R:11-18-25	W:26-23
13.R:5-14(f)	W:23-19
14.R:25-29	W:31-26
15.R:6-10	W:19-16
16.R:8-12	W:16-11
17.R:12-16	W:28-24
18.R:29-25	W:32-27
19.R:16-20	W:24-19
20.R:13-17	W:26-23

21.R:25-22 W:19-16 22.R:22-26 W:23-19 23.R:26-31 W:27-23 24.R:17-22 W:11-7
25.R:2-11(f) W:16-7(f) 26.R:31-27 W:7-2 27.R:27-18(f) W:19-16 28.R:18-23 W:16-11
29.R:22-26 W:2-6

[probably a mistake by human]

Red

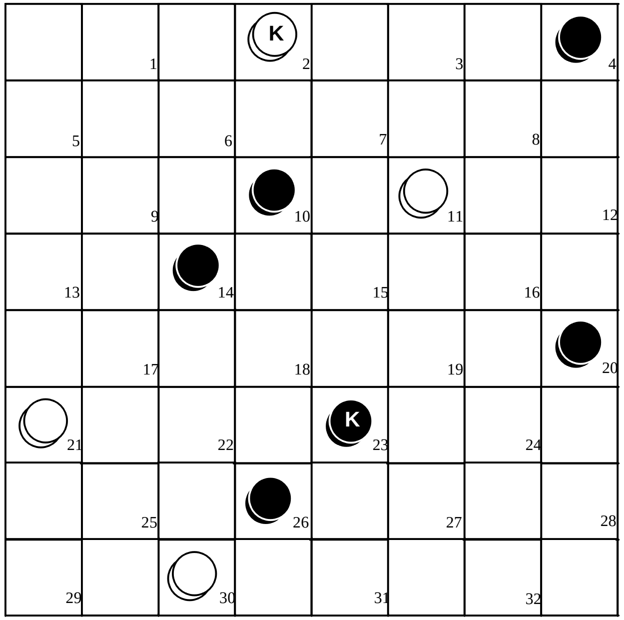

White

Game against Human rated 1771: Win

ENN	Human
29.R:4-8,	W:11-4
30.R:20-24,	W:6-15(f)
31.R:23-27,	W:30-23(f)
32.R:27-18-11	W:21-17 [figure]
33.R:14-21(f)	W:4-8(f)
34.R:11-4(f)	

[game over, red (computer) wins]

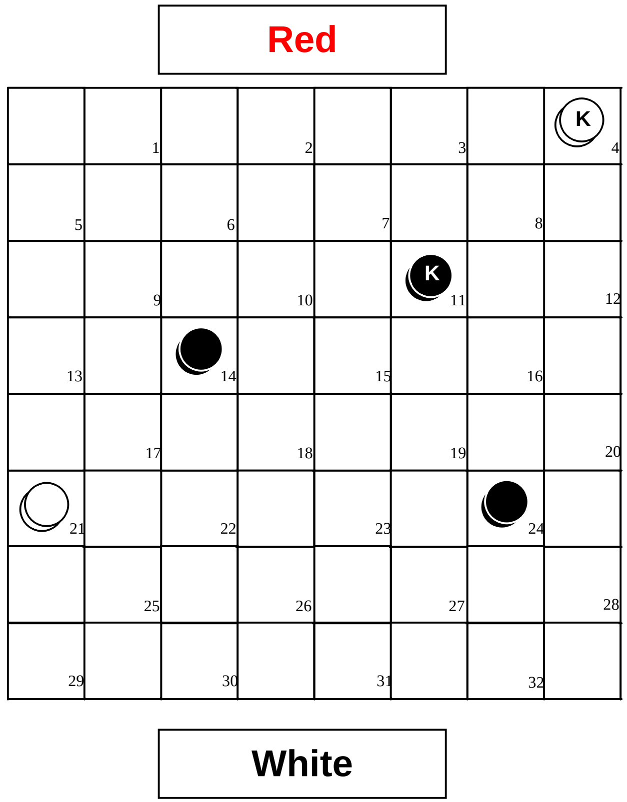

Extension to Spatial Neural Network

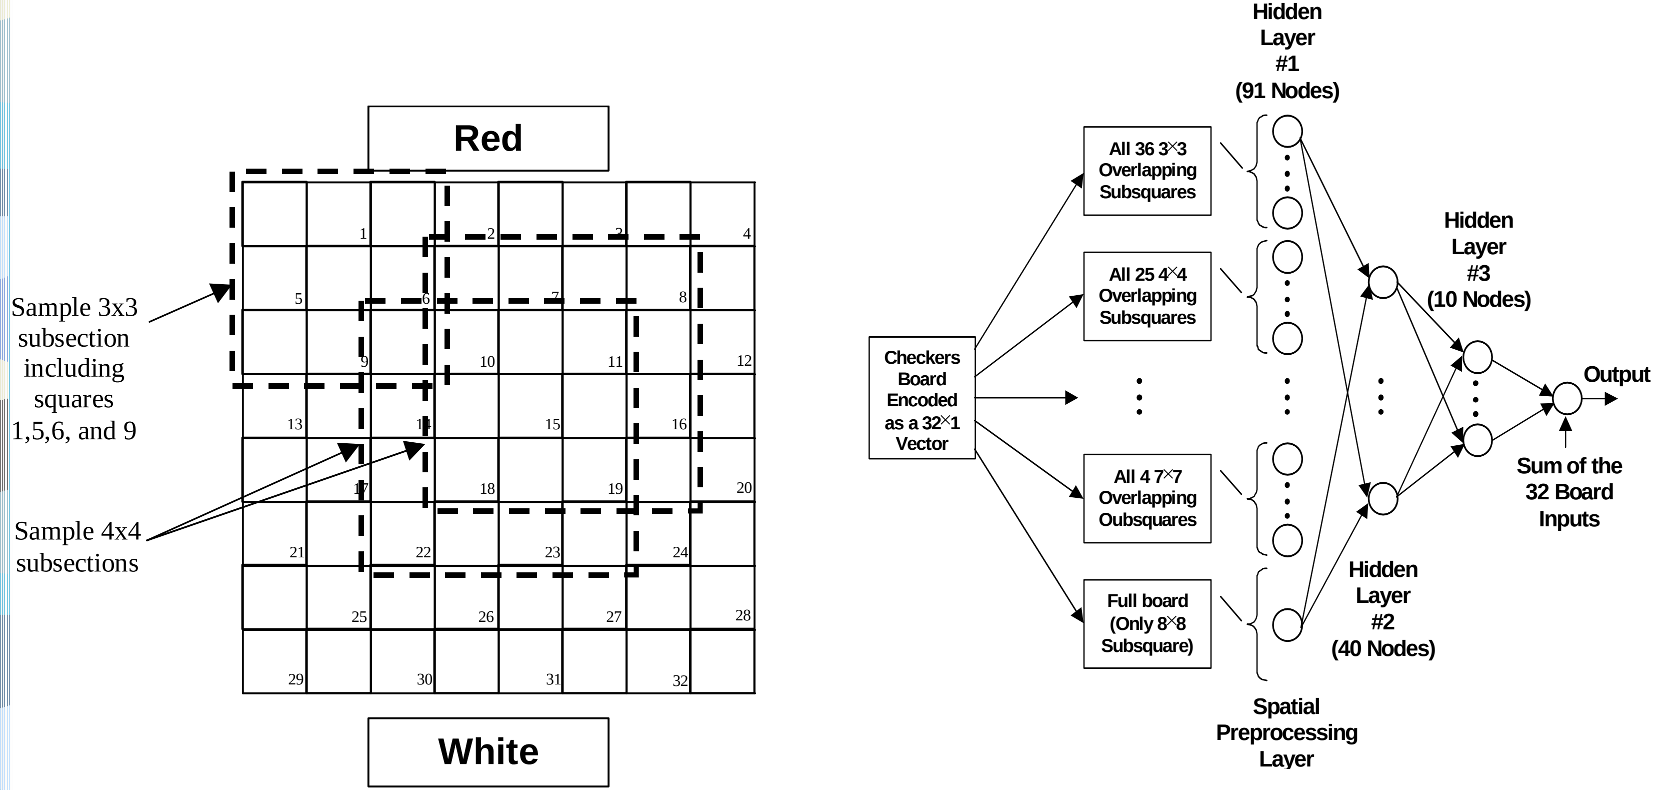

Results with Spatial Neural Network

- Trained over 840 generations
- Tested on 150+ games
- Rated at 2040, “expert” level
- Often played to restrict mobility of opponent
  - To the extent that the neural network used this feature, it first *had to invent the feature*
  - We named the neural network Anaconda

Hoyle® Classic Games

- You play against characters
- Each has a different skill level
- A six-game match was played against characters of "expert" ability
- Anaconda won 6-0

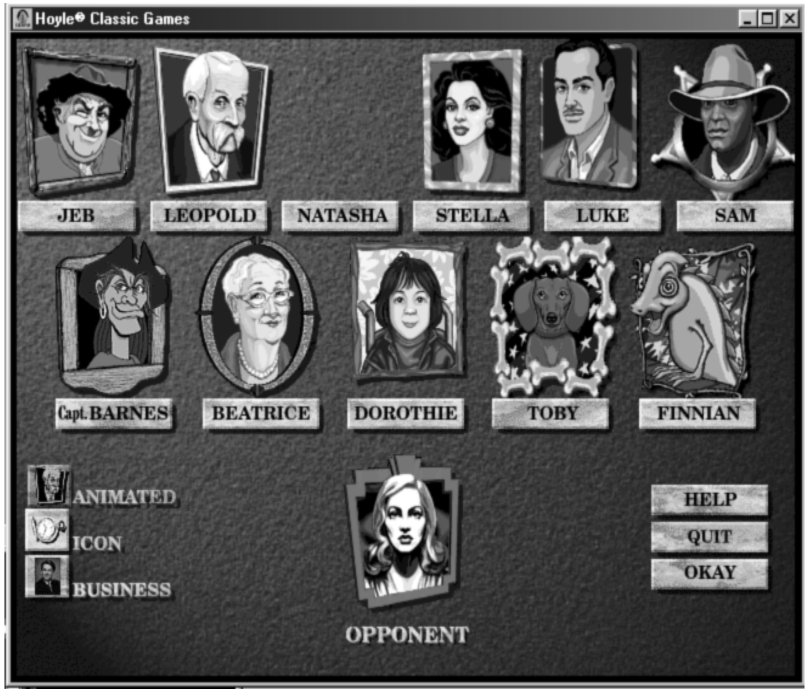

What Has Been Accomplished?

- Co-evolution of checkers players using
  - A representation defining the location of pieces on the board
  - A variable coding value for the king
  - A heuristic (DFS) for searching ahead 6-ply
  - A heuristic (minimax) for selecting which move to favor in light of the NN evaluation function
  - The potential to use piece differential as a feature
- No expert knowledge
- The evolved player achieved a rating of 1750.8 at www.zone.com, placing the ENN as a better than median Class B player.
- Anaconda has taught itself to play like an expert

What's Next?

- "Darwinian invention and problem solving" are within our grasp
- How far are we from evolving "brains" that are comparable with humans?
  - 2 years?
  - 20 Years?
  - Never?
- Important to remain focused on the mechanisms that underlie the "intelligent" behavior otherwise we have every opportunity to make the answer "never"